CS252
Graduate Computer Architecture
Lecture 10

Prediction/Speculation
February 22nd, 2012

John Kubiatowicz
Electrical Engineering and Computer Sciences
University of California, Berkeley

http://www.eecs.berkeley.edu/~kubitron/cs252

Review: Independent "Fetch" unit

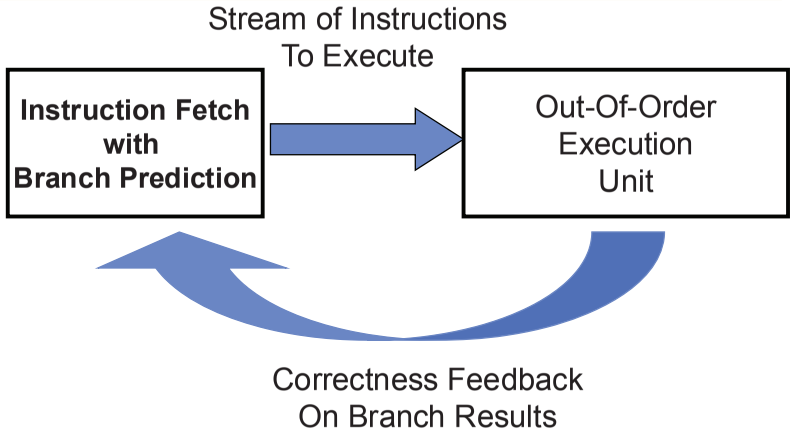

- **Instruction fetch decoupled from execution**
- **Often issue logic (+ rename) included with Fetch**

Review:Dynamic Branch Prediction Problem

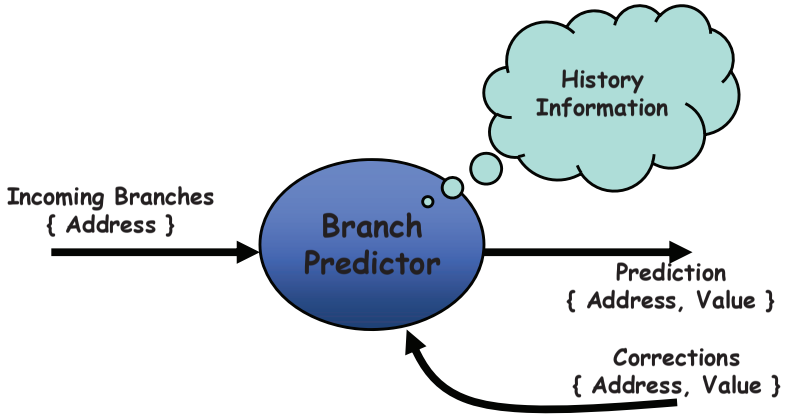

- **Incoming stream of addresses**
- **Fast outgoing stream of predictions**
- **Correction information returned from pipeline**

Branch Target Buffer

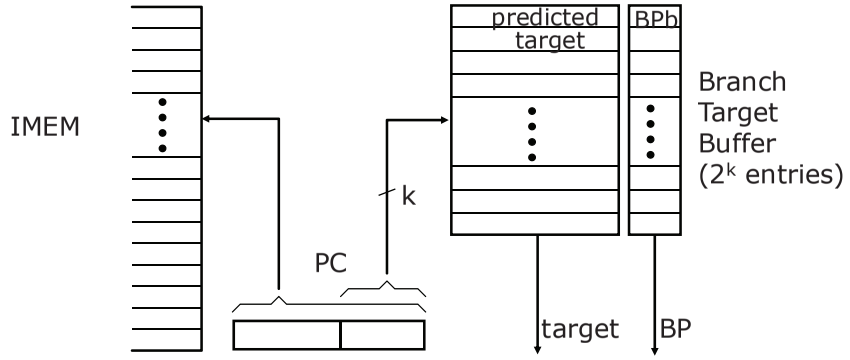

BP bits are stored with the predicted target address.

IF stage: *If (BP=taken) then nPC=target else nPC=PC+4*
later: *check prediction, if wrong then kill the instruction and update BTB & BPb else update BPb*

Combining BTB and BHT

- BTB entries are considerably more expensive than BHT, but can redirect fetches at earlier stage in pipeline and can accelerate indirect branches (JR)
- BHT can hold many more entries and is more accurate

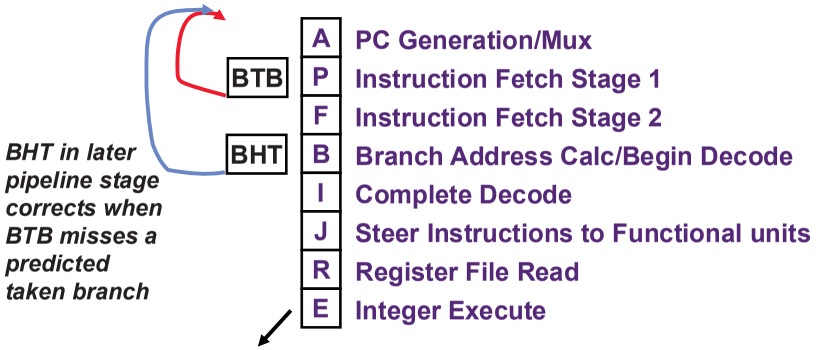

BTB/BHT only updated after branch resolves in E stage

Correlating Branches

- Hypothesis: recent branches are correlated; that is, behavior of recently executed branches affects prediction of current branch
- Two possibilities; Current branch depends on:
  - Last m most recently executed branches anywhere in program Produces a "GA" (for "global adaptive") in the Yeh and Patt classification (e.g. GAg)
  - Last m most recent outcomes of same branch. Produces a "PA" (for "per-address adaptive") in same classification (e.g. PAg)
- Idea: record m most recently executed branches as taken or not taken, and use that pattern to select the proper branch history table entry
  - A single history table shared by all branches (appends a "g" at end), indexed by history value.
  - Address is used along with history to select table entry (appends a "p" at end of classification)
  - If only portion of address used, often appends an "s" to indicate "set-indexed" tables (I.e. GAs)

Exploiting Spatial Correlation

Yeh and Patt, 1992

```
if (x[i] < 7) then
      y += 1;
if (x[i] < 5) then
      c -= 4;
```

If first condition false, second condition also false

History register, H, records the direction of the last N branches executed by the processor

Two-Level Branch Predictor (e.g. GAs)

Pentium Pro uses the result from the last two branches to select one of the four sets of BHT bits (~95% correct)

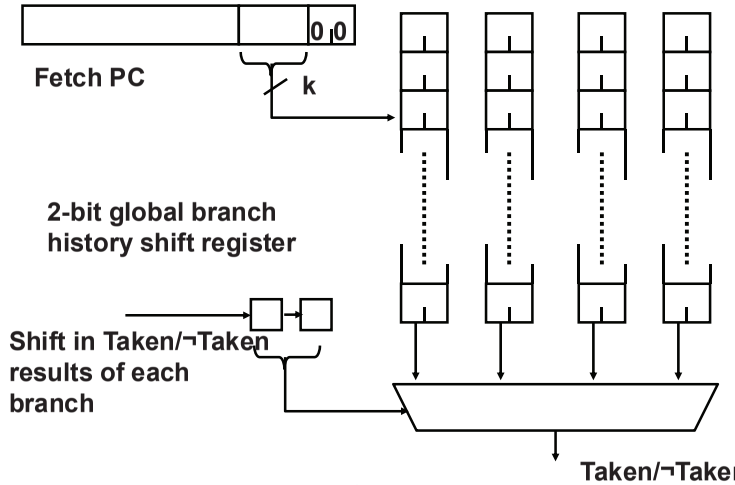

What are Important Metrics?

- Clearly, Hit Rate matters
  - Even 1% can be important when above 90% hit rate
- Speed: Does this affect cycle time?
- Space: Clearly Total Space matters!
  - Papers which do not try to normalize across different options are playing fast and lose with data
  - Try to get best performance for the cost

Accuracy of Different Schemes

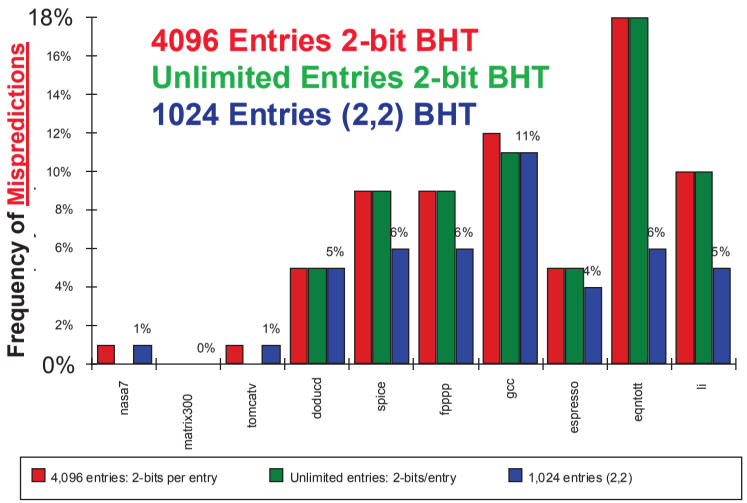

BHT Accuracy

- Mispredict because either:
  - Wrong guess for that branch
  - Got branch history of wrong branch when index the table
- 4096 entry table programs vary from 1% misprediction (nasa7, tomcatv) to 18% (eqntott), with spice at 9% and gcc at 12%
  - For SPEC92, 4096 about as good as infinite table
- How could HW predict "this loop will execute 3 times" using a simple mechanism?
  - Need to track history of just that branch
  - For given pattern, track most likely following branch direction
- Leads to two separate types of recent history tracking:
  - GBHR (Global Branch History Register)
  - PABHR (Per Address Branch History Table)
- Two separate types of Pattern tracking
  - GPHT (Global Pattern History Table)
  - PAPHT (Per Address Pattern History Table)

Yeh and Patt classification

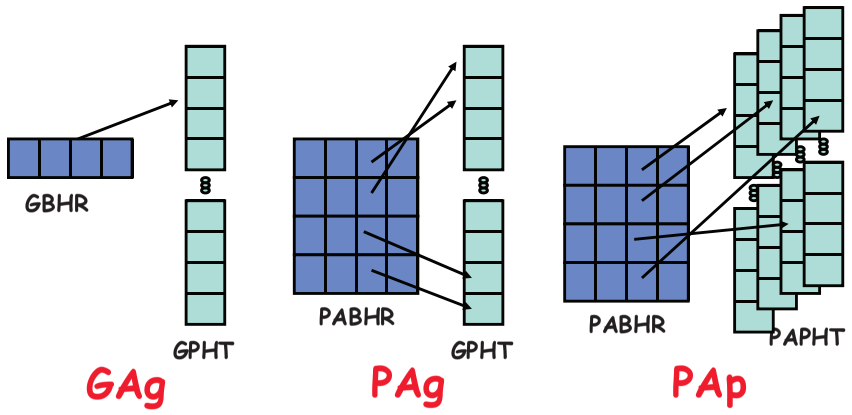

- GAg: Global History Register, Global History Table
- PAg: Per-Address History Register, Global History Table
- PAp: Per-Address History Register, Per-Address History Table

Two-Level Adaptive Schemes: History Registers of Same Length (6 bits)

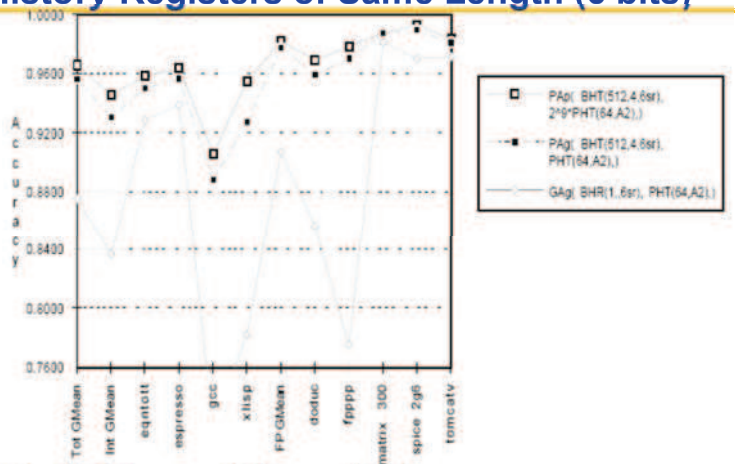

- PAp best: But uses a lot more state!
- GAg not effective with 6-bit history registers
  - Every branch updates the same history register⇒interference
- PAg performs better because it has a branch history table

Versions with Roughly same accuracy (97%)

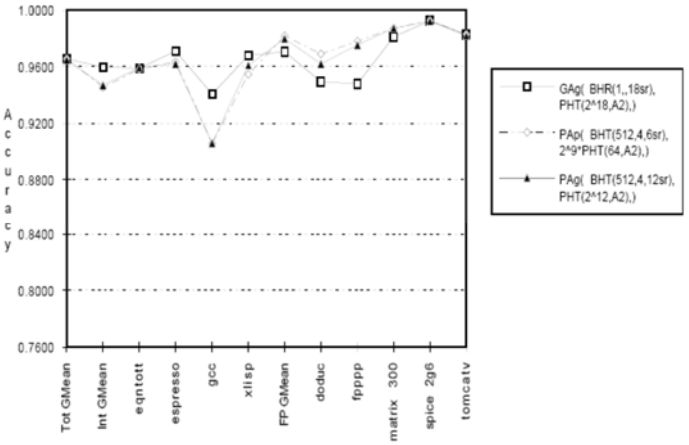

- Cost:
  - GAg requires 18-bit history register
  - PAg requires 12-bit history register
  - PAp requires 6-bit history register
- PAg is the cheapest among these

Why doesn't GAg do better?

- Difference between GAg and both PA variants:
  - GAg tracks correllations between different branches
  - PAg/PAp track corellations between different instances of the same branch
- These are two different types of pattern tracking
  - Among other things, GAg good for branches in straight-line code, while PA variants good for loops
- Problem with GAg? It aliases results from different branches into same table
  - Issue is that different branches may take same global pattern and resolve it differently
  - GAg doesn't leave flexibility to do this

Other Global Variants: Try to Avoid Aliasing

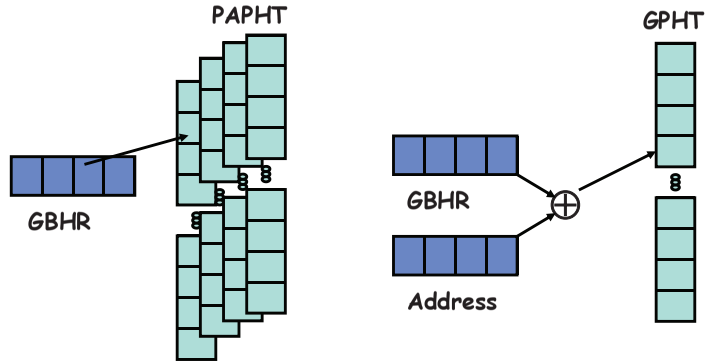

- GAs: Global History Register, Per-Address (Set Associative) History Table
- Gshare: Global History Register, Global History Table with Simple attempt at anti-aliasing

Branches are Highly Biased

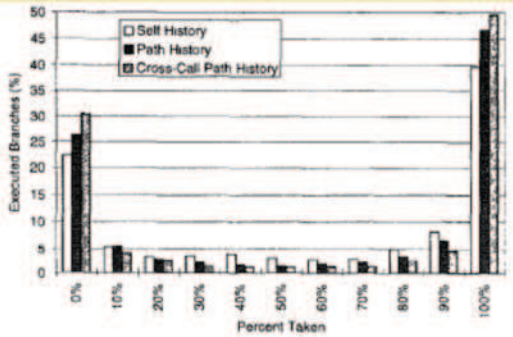

- From: "A Comparative Analysis of Schemes for Correlated Branch Prediction," by Cliff Young, Nicolas Gloy, and Michael D. Smith
- Many branches are highly biased to be taken or not taken
  - Use of path history can be used to further bias branch behavior
- Can we exploit bias to better predict the unbiased branches?
  - Yes: filter out biased branches to save prediction resources for the unbiased ones

Exploiting Bias to avoid Aliasing: Bimode and YAGS

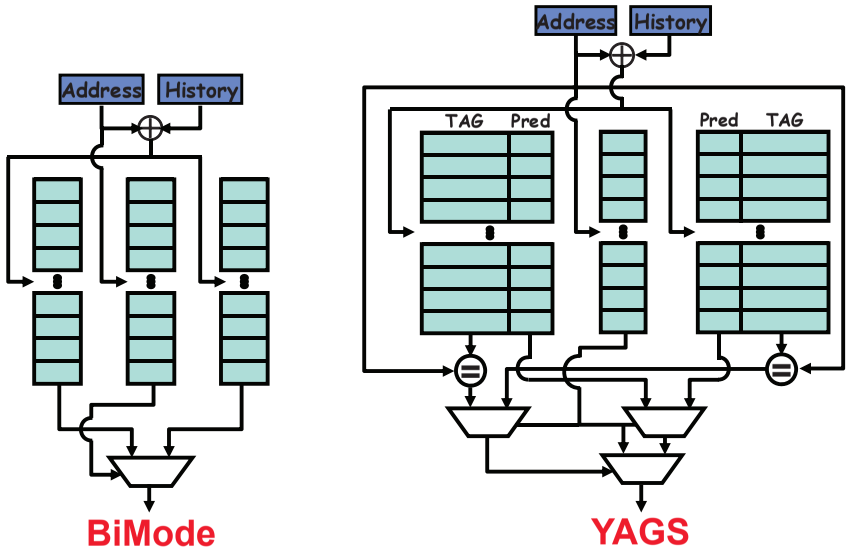

Is Global or Local better?

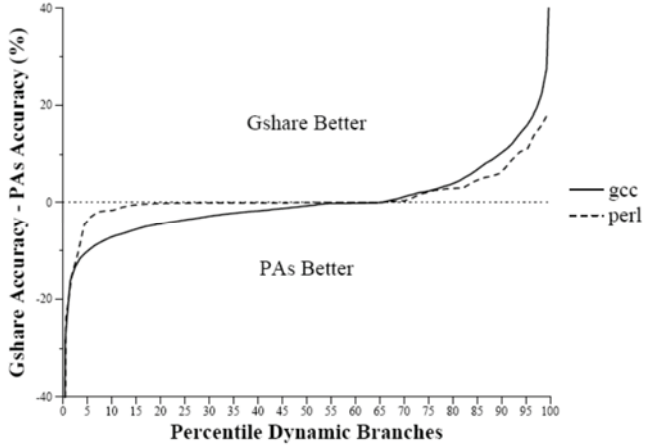

- Neither: Some branches local, some global
  - From: "An Analysis of Correlation and Predictability: What Makes Two-Level Branch Predictors Work," Evers, Patel, Chappell, Patt
  - Difference in predictability quite significant for some branches!

Dynamically finding structure in Spaghetti

- Consider complex "spaghetti code"
- Are all branches likely to need the same type of branch prediction?
  - No.
- What to do about it?
  - How about predicting which predictor will be best?
  - Called a "Tournament predictor"

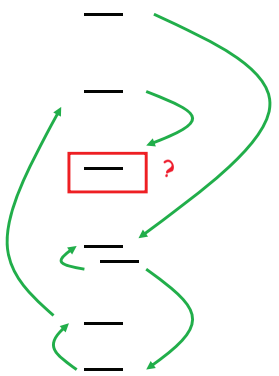

Tournament Predictors

- Motivation for correlating branch predictors is 2-bit predictor failed on important branches; by adding global information, performance improved
- Tournament predictors: use 2 predictors, 1 based on global information and 1 based on local information, and combine with a selector
- Use the predictor that tends to guess correctly

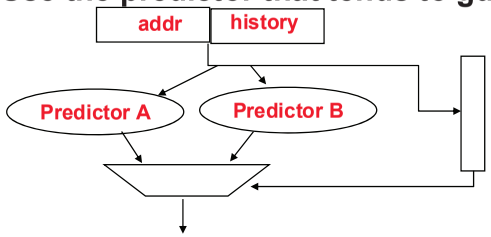

Tournament Predictor in Alpha 21264

- 4K 2-bit counters to choose from among a global predictor and a local predictor
- Global predictor also has 4K entries and is indexed by the history of the last 12 branches; each entry in the global predictor is a standard 2-bit predictor
  - 12-bit pattern: ith bit 0 => ith prior branch not taken; ith bit 1 => ith prior branch taken;
- Local predictor consists of a 2-level predictor:
  - Top level a local history table consisting of 1024 10-bit entries; each 10-bit entry corresponds to the most recent 10 branch outcomes for the entry. 10-bit history allows patterns 10 branches to be discovered and predicted.
  - Next level Selected entry from the local history table is used to index a table of 1K entries consisting a 3-bit saturating counters, which provide the local prediction
- Total size: 4K*2 + 4K*2 + 1K*10 + 1K*3 = 29K bits!
  (~180,000 transistors)

% of predictions from local predictor in Tournament Scheme

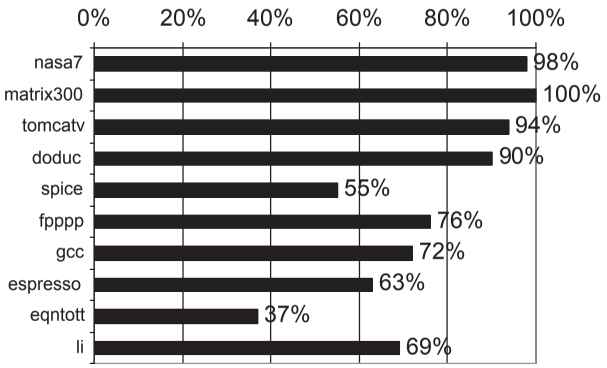

Accuracy of Branch Prediction

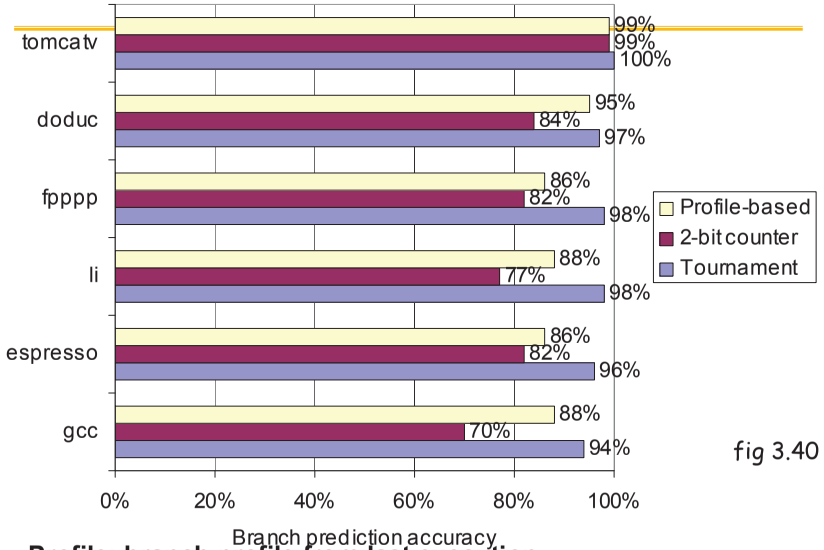

fig 3.40

- Profile: branch profile from last execution (static in that in encoded in instruction, but profile)

Accuracy v. Size (SPEC89)

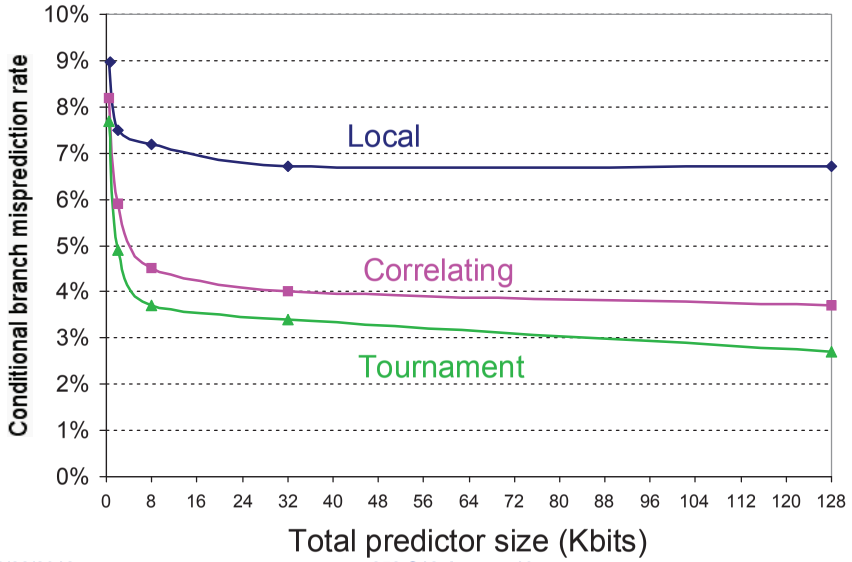

Pitfall: Sometimes bigger and dumber is better

- 21264 uses tournament predictor (29 Kbits)
- Earlier 21164 uses a simple 2-bit predictor with 2K entries (or a total of 4 Kbits)
- SPEC95 benchmarks, 21264 outperforms
  - 21264 avg. 11.5 mispredictions per 1000 instructions
  - 21164 avg. 16.5 mispredictions per 1000 instructions
- Reversed for transaction processing (TP) !
  - 21264 avg. 17 mispredictions per 1000 instructions
  - 21164 avg. 15 mispredictions per 1000 instructions
- TP code much larger & 21164 hold 2X branch predictions based on local behavior (2K vs. 1K local predictor in the 21264)

Administrivia

- Midterm I: Wednesday 3/21
  Location: 405 Soda Hall
  TIME: 5:00-8:00
  - Can have 1 sheet of 8½x11 *handwritten* notes – both sides
  - No microfiche of the book!
- Meet at LaVal's afterwards for Pizza and Beverages
  - Great way for me to get to know you better
  - I'll Buy!
- CS252 First Project proposal due by Friday 3/2
  - Need two people/project (although can justify three for right project)
  - Complete Research project in 9 weeks
    - » Typically investigate hypothesis by building an artifact and measuring it against a "base case"
    - » Generate conference-length paper/give oral presentation
    - » Often, can lead to an actual publication.

One important tool is RAMP Gold: FAST Emulation of new Hardware

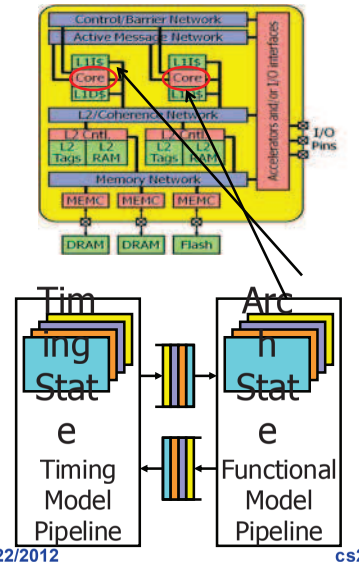

- RAMP emulation model for Parlab manycore
  - SPARC v8 ISA -> v9
  - Considering ARM model
- Single-socket manycore target
- Split functional/timing model, both in hardware
  - Functional model: Executes ISA
  - Timing model: Capture pipeline timing detail (can be cycle accurate)
- Host multithreading of both functional and timing models
- Built for Virtex-5 systems (ML505 or BEE3)
- Have Tessellation OS currently running on RAMP system!

Tessellation: The Exploded OS

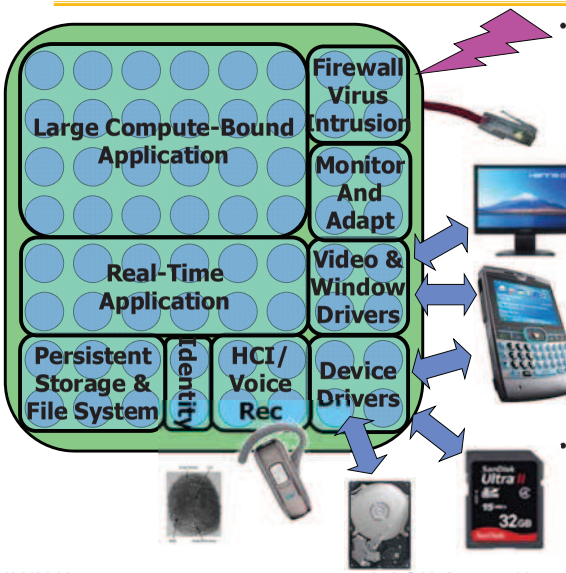

- Normal Components split into pieces
  - Device drivers (Security/Reliability)
  - Network Services (Performance)
    - » TCP/IP stack
    - » Firewall
    - » Virus Checking
    - » Intrusion Detection
  - Persistent Storage (Performance, Security, Reliability)
  - Monitoring services
    - » Performance counters
    - » Introspection
  - Identity/Environment services (Security)
    - » Biometric, GPS, Possession Tracking
- Applications Given Larger Partitions
  - Freedom to use resources arbitrarily

Implementing the Space-Time Graph

- Partition Policy layer (allocation)
  - Allocates Resources to Cells based on Global policies
  - Produces only implementable space-time resource graphs
  - May deny resources to a cell that requests them (admission control)
- Mapping layer (distribution)
  - Makes no decisions
  - Time-Slices at a course granularity (when time-slicing necessary)
  - performs bin-packing like operation to implement space-time graph
  - In limit of *many* processors, no time multiplexing processors, merely distributing resources
- Partition Mechanism Layer
  - Implements hardware partitions and secure channels
  - Device Dependent: Makes use of more or less hardware support for QoS and Partitions

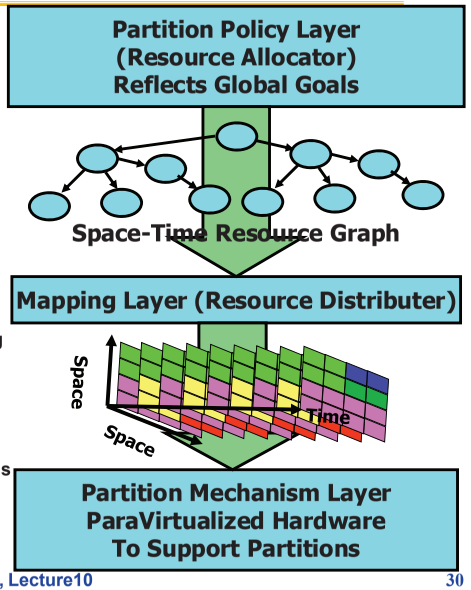

Sample of what could make good projects

- Implement new resource partitioning mechanisms on RAMP and integrate into OS
  - You can actually develop a new hardware mechanism, put into the OS, and show how partitioning gives better performance or real-time behavior
  - You could develop new message-passing interfaces and do the same
- Virtual I/O devices
  - RAMP-Gold runs in virtual time
  - Develop devices and methodology for investigating real-time behavior of these devices in Tessellation running on RAMP
- Energy monitoring and adaptation
  - How to measure energy consumed by applications and adapt accordingly
  - We have access to energy counters in some hardware
- Develop and evaluate new parallel communication model
  - Target for Multicore systems
  - New Message-Passing Interface, New Network Routing Layer
- Investigate applications under different types of hardware
  - CUDA vs MultiCore, etc
- New Style of computation, tweak on existing one
- Better Memory System, etc

Projects continued

- Optimize a parallel algorithm for on-chip message passing (requires access to FPGA and/or Tilera chip)
- New hardware support for security
  - Is there some universal primitives that could be added to multicore chips to make sure that data stays private/secure?
- Develop an energy model for the RAMP Gold
- How do different cache-coherence protocols affect a chip's energy-efficiency?
  - Propose an optimization(s) to the CC protocol.-
- Propose "the next FPU" - the next fixed-function accelerator with a wide range of applicability.
  - Simulate (chisel, verilog) and compare performance & energy versus doing it in software. What's the interface?

Projects using Quantum CAD Flow

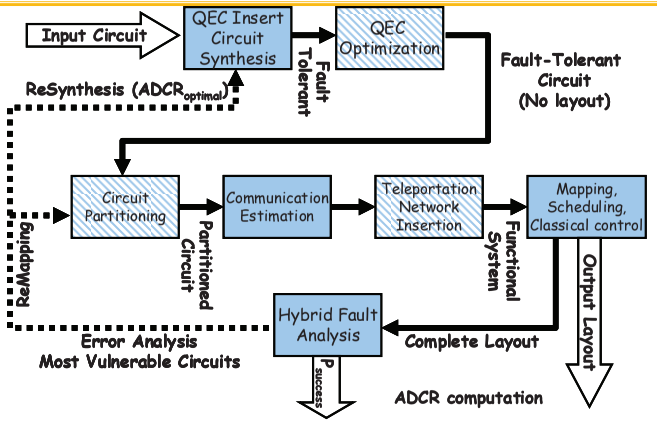

- **Use the quantum CAD flow developed in Kubiatowicz's group to investigate Quantum Circuits**
  - Tradeoff in area vs performance for Shor's algorithm
  - Other interesting algorithms (Quantum Simulation)

Speculating Both Directions

An alternative to branch prediction is to execute both directions of a branch *speculatively*

- resource requirement is proportional to the number of concurrent speculative executions
- only half the resources engage in useful work when both directions of a branch are executed speculatively
- branch prediction takes less resources than speculative execution of both paths

With accurate branch prediction, it is more cost effective to dedicate all resources to the predicted direction

Review: Memory Disambiguation

- **Question: Given a load that follows a store in program order, are the two related?**
  - Trying to detect RAW hazards through memory
  - Stores commit in order (ROB), so no WAR/WAW memory hazards.
- **Implementation**
  - Keep queue of stores, in program order
  - Watch for position of new loads relative to existing stores
  - Typically, this is a different buffer than ROB!
    - » Could be ROB (has right properties), but too expensive
- **When have address for load, check store queue:**
  - If *any* store prior to load is waiting for its address⇒?????
  - If load address matches earlier store address (associative lookup), then we have a *memory-induced RAW hazard*:
    - » store value available ⇒ return value
    - » store value not available ⇒ return ROB number of source
  - Otherwise, send out request to memory

In-Order Memory Queue

- **Execute all loads and stores in program order**

=> Load and store cannot leave ROB for execution until all previous loads and stores have completed execution

- **Can still execute loads and stores speculatively, and out-of-order with respect to other instructions**

Conservative O-o-O Load Execution

```
st r1, (r2)
ld r3, (r4)
```

- Split execution of store instruction into two phases: address calculation and data write
- Can execute load before store, if addresses known and r4 != r2
- Each load address compared with addresses of all previous uncommitted stores *(can use partial conservative check i.e., bottom 12 bits of address)*
- Don't execute load if any previous store address not known

(MIPS R10K, 16 entry address queue)

Address Speculation

```
st r1, (r2)
ld r3, (r4)
```

- Guess that r4 != r2
- Execute load before store address known
- Need to hold all completed but uncommitted load/store addresses in program order
- If subsequently find r4==r2, squash load and *all* following instructions

=> Large penalty for inaccurate address speculation

Memory Dependence Prediction
(Alpha 21264)

```
st r1, (r2)
ld r3, (r4)
```

- Guess that r4 != r2 and execute load before store
- If later find r4==r2, squash load and all following instructions, but mark load instruction as *store-wait*
- Subsequent executions of the same load instruction will wait for all previous stores to complete
- Periodically clear *store-wait* bits

Speculative Store Buffer

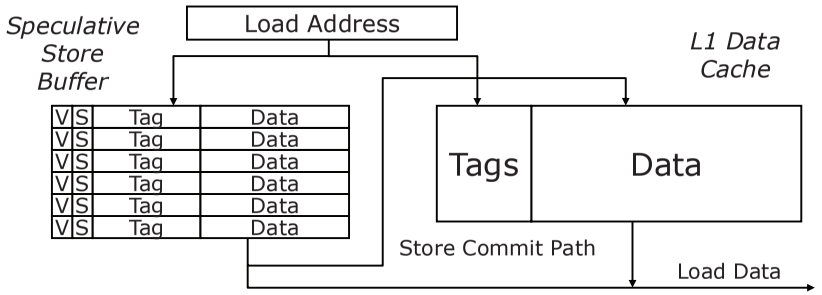

- On store execute:
  - mark entry valid and speculative, and save data and tag of instruction.
- On store commit:
  - clear speculative bit and eventually move data to cache
- On store abort:
  - clear valid bit

Speculative Store Buffer

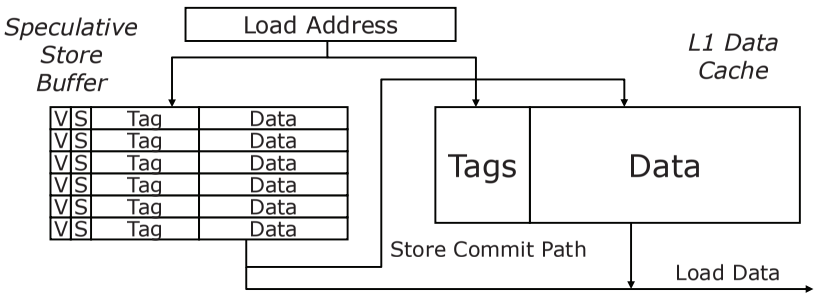

- If data in both store buffer and cache, which should we use:
  Speculative store buffer
- If same address in store buffer twice, which should we use:
  Youngest store older than load

Memory Dependence Prediction

- Important to speculate? Two Extremes:
  - Naïve Speculation: always let load go forward
  - No Speculation: always wait for dependencies to be resolved
- Compare Naïve Speculation to No Speculation
  - False Dependency: wait when don't have to
  - Order Violation: result of speculating incorrectly
- Goal of prediction:
  - Avoid false dependencies *and* order violations

	Naive Speculation		No Speculation
Spec95 Program	Memory Order Viols Per 1K Instrs	Memory Trap Penalty (Cycles)	False Dep. Per 1K Instrs
go	6	13	157
m88ksim	20	12	168
gcc	5	15	187
compress	11	15	129
xlisp	11	14	179
ijpeg	23	15	150
perl prim	20	15	215
perl scrab	10	15	185
vortex	7	19	215
tomcatv	4	22	264
swim	2	36	224
mgrid	0	18	262
applu	18	22	212
apsi	7	35	247
fpppp	10	17	275
wave5	24	21	188
turb3d	6	16	213

From "Memory Dependence Prediction using Store Sets", Chrysos and Emer.

Said another way: Could we do better?

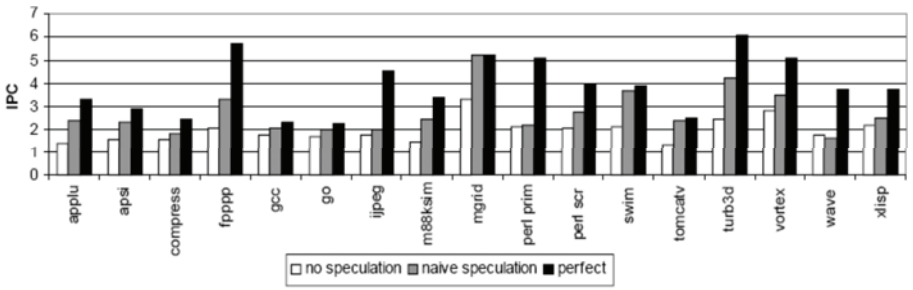

- Results from same paper: performance improvement with oracle predictor
  - We can get significantly better performance if we find a good predictor
  - Question: How to build a good predictor?

Premise: Past indicates Future

- Basic Premise is that past dependencies indicate future dependencies
  - Not always true! Hopefully true most of time
- Store Set: Set of store insts that affect given load
  - Example:

Addr	Inst
0	Store C
4	Store A
8	Store B
12	Store C
28	Load B ⇒ Store set { PC 8 }
32	Load D ⇒ Store set { (null) }
36	Load C ⇒ Store set { PC 0, PC 12 }
40	Load B ⇒ Store set { PC 8 }

  - Idea: Store set for load starts empty. If ever load go forward and this causes a violation, add offending store to load's store set
- Approach: For each indeterminate load:
  - If Store from Store set is in pipeline, stall
    Else let go forward
- Does this work?

How well does an infinite tracking work?

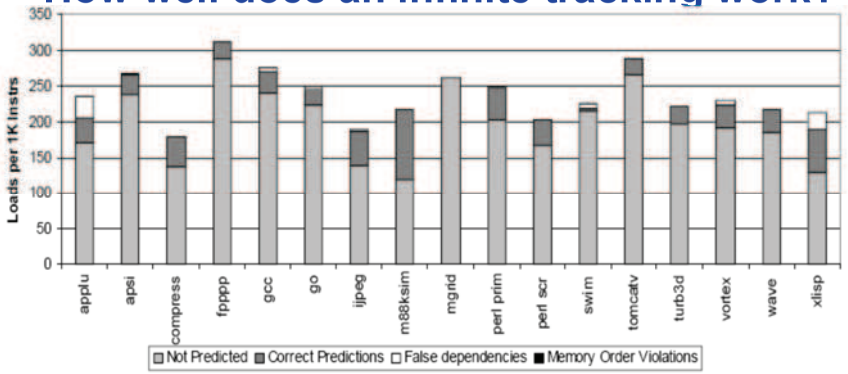

- "Infinite" here means to place no limits on:
  - Number of store sets
  - Number of stores in given set
- Seems to do pretty well
  - Note: "Not Predicted" means load had empty store set
  - Only Applu and Xlisp seems to have false dependencies

How to track Store Sets in reality?

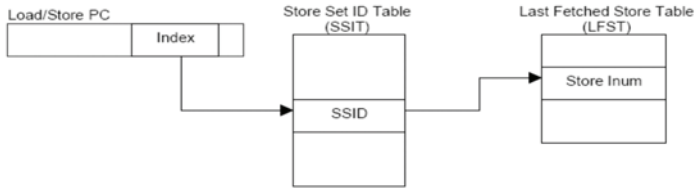

- SSIT: Assigns Loads and Stores to Store Set ID (SSID)
  - Notice that this requires each store to be in only one store set!
- LFST: Maps SSIDs to most recent fetched store
  - When Load is fetched, allows it to find most recent store in its store set that is executing (if any) ⇒ allows stalling until store finished
  - When Store is fetched, allows it to wait for previous store in store set
    - » Pretty much same type of ordering as enforced by ROB anyway
    - » Transitivity⇒ loads end up waiting for all active stores in store set
- What if store needs to be in two store sets?
  - Allow store sets to be merged together deterministically
    - » Two loads, multiple stores get same SSID
- Want periodic clearing of SSIT to avoid:
  - problems with aliasing across program
  - Out of control merging

How well does this do?

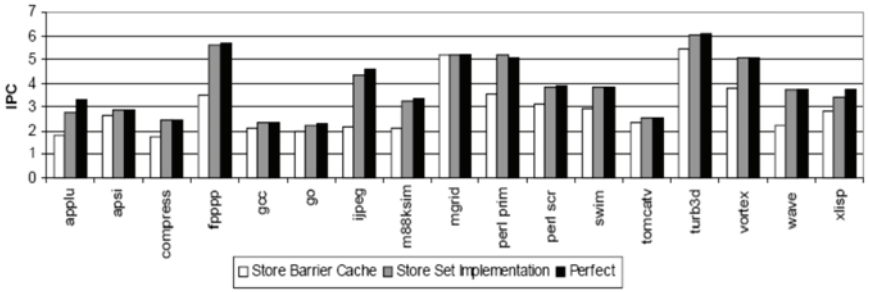

- Comparison against Store Barrier Cache
  - Marks individual Stores as "tending to cause memory violations"
  - Not specific to particular loads….
- Problem with APPLU?
  - Analyzed in paper: has complex 3-level inner loop in which loads occasionally depend on stores
  - Forces overly conservative stalls (i.e. false dependencies)

Load Value Predictability

- Try to predict the result of a load before going to memory
- Paper: "Value locality and load value prediction"
  - Mikko H. Lipasti, Christopher B. Wilkerson and John Paul Shen
- Notion of value locality
  - Fraction of instances of a given load that match last n different values
- Is there any value locality in typical programs?
  - Yes!
  - With history depth of 1: most integer programs show over 50% repetition
  - With history depth of 16: most integer programs show over 80% repetition
  - Not everything does well: see *cjpeg, swm256, and tomcatv*
- Locality varies by type:
  - Quite high for inst/data *addresses*
  - Reasonable for integer values
  - Not as high for FP values

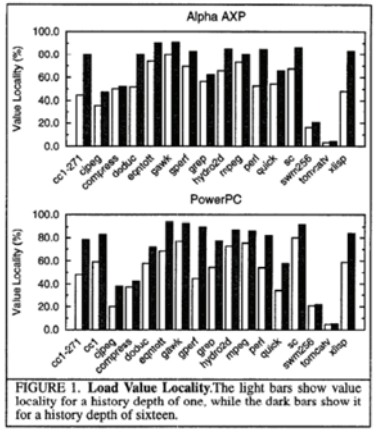

FIGURE 1. **Load Value Locality.** The light bars show value locality for a history depth of one, while the dark bars show it for a history depth of sixteen.

Load Value Prediction Table

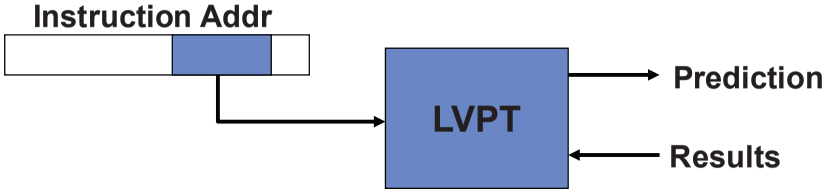

- Load Value Prediction Table (LVPT)
  - Untagged, Direct Mapped
  - Takes Instructions ⇒ Predicted Data
- Contains *history* of last n unique values from given instruction
  - Can contain aliases, since untagged
- How to predict?
  - When n=1, easy
  - When n=16? Use Oracle
- Is every load predictable?
  - No! Why not?
  - Must identify predictable loads somehow

Load Classification Table (LCT)

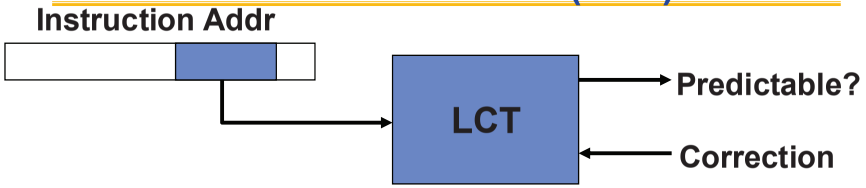

- Load Classification Table (LCT)
  - Untagged, Direct Mapped
  - Takes Instructions ⇒ Single bit of whether or not to predict
- How to implement?
  - Uses saturating counters (2 or 1 bit)
  - When prediction correct, increment
  - When prediction incorrect, decrement
- With 2 bit counter
  - 0,1 ⇒ not predictable
  - 2 ⇒ predictable
  - 3 ⇒ constant (very predictable)
- With 1 bit counter
  - 0 ⇒ not predictable
  - 1 ⇒ constant (very predictable)

Accuracy of LCT

- Question of accuracy is about how well we avoid:
  - Predicting unpredictable load
  - Not predicting predictable loads
- How well does this work?
  - Difference between "Simple" and "Limit": history depth
    - » Simple: depth 1
    - » Limit: depth 16
  - Limit tends to classify more thing as predictable (since this works more often)
- Basic Principle:
  - Often works better to have one structure decide on the basic "predictability" of structure
  - Independent of prediction structure

Bench-mark	PowerPC				Alpha AXP			
	Simple		Limit		Simple		Limit	
	Unpr	Pred	Unpr	Pred	Unpr	Pred	Unpr	Pred
cc1-271	86%	64%	58%	90%	86%	57%	64%	86%
cjpeg	97%	61%	92%	61%	93%	75%	93%	82%
compress	99%	94%	97%	90%	98%	56%	97%	94%
doduc	83%	75%	82%	92%	84%	68%	78%	92%
eqntott	91%	85%	88%	99%	68%	80%	83%	97%
gawk	85%	92%	44%	95%	74%	86%	59%	93%
gperf	93%	75%	76%	97%	77%	79%	77%	91%
grep	93%	88%	67%	81%	85%	82%	92%	92%
hydro2d	82%	85%	63%	91%	86%	80%	60%	89%
mpeg	86%	90%	78%	93%	84%	88%	85%	93%
perl	84%	71%	65%	93%	83%	66%	74%	93%
quick	98%	84%	93%	89%	98%	95%	96%	95%
sc	77%	90%	59%	97%	86%	85%	78%	95%
swm256	99%	89%	99%	93%	99%	86%	99%	90%
tomcatv	100%	89%	100%	98%	99%	68%	99%	70%
xlisp	88%	83%	77%	93%	90%	74%	76%	93%
GM	90%	81%	75%	90%	86%	78%	81%	90%

Constant Value Unit

- Idea: Identify a load instruction as "constant"
  - Can ignore cache lookup (no verification)
  - Must enforce by monitoring result of stores to remove "constant" status
- How well does this work?
  - Seems to identify 6-18% of loads as constant
  - Must be unchanging enough to cause LCT to classify as constant

Benchmark	PowerPC		Alpha AXP	
	Simple	Limit	Simple	Limit
cc1-271	13%	23%	10%	14%
cjpeg	4%	7%	17%	17%
compress	33%	34%	36%	42%
doduc	5%	20%	5%	15%
eqntott	19%	44%	21%	35%
gawk	10%	28%	31%	31%
gperf	21%	39%	38%	56%
grep	16%	24%	18%	22%
hydro2d	2%	8%	3%	10%
mpeg	12%	25%	10%	28%
perl	8%	19%	7%	8%
quick	0%	0%	31%	31%
sc	32%	46%	26%	31%
swm256	8%	17%	12%	12%
tomcatv	0%	0%	1%	1%
xlisp	14%	45%	8%	30%
GM	6%	11%	12%	18%

Load Value Architecture

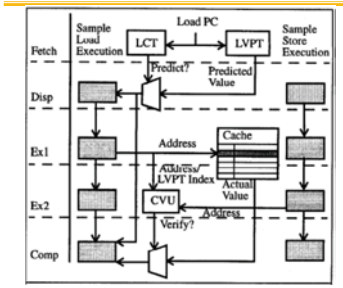

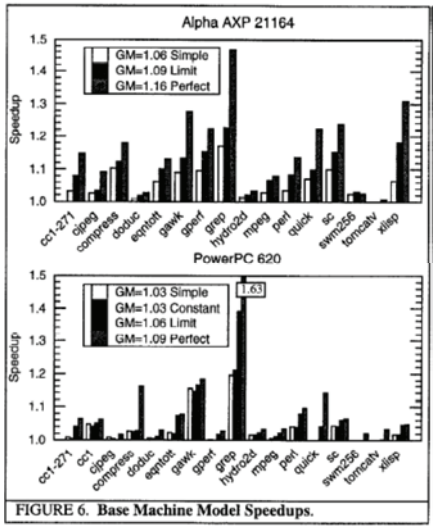

FIGURE 6. Base Machine Model Speedups.

- LCT/LVPT in fetch stage
- CVU in execute stage
  - Used to bypass cache entirely
  - (Know that result is good)
- Results: Some speedups
  - 21264 seems to do better than Power PC
  - Authors think this is because of small first-level cache and in-order execution makes CVU more useful

Data Value Prediction

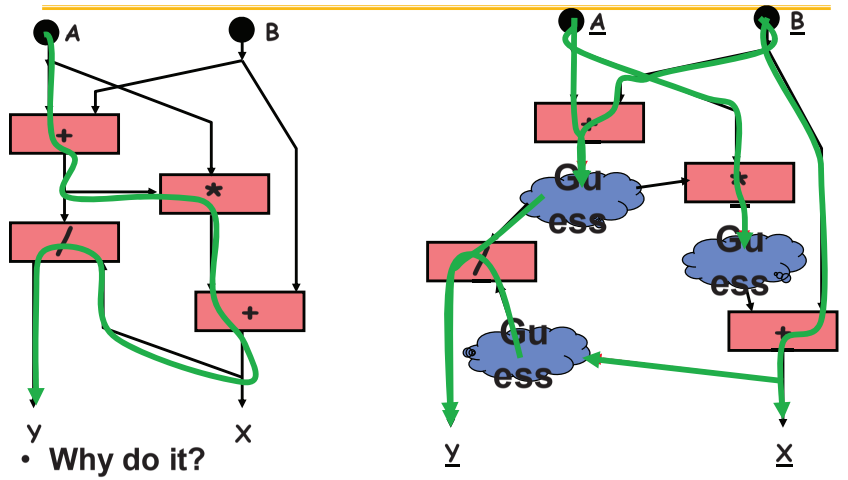

- Why do it?
  - Can "Break the DataFlow Boundary"
  - Before: Critical path = 4 operations (probably worse)
  - After: Critical path = 1 operation (plus verification)

Data Value Predictability

- "The Predictability of Data Values"
  - Yiannakis Sazeides and James Smith, Micro 30, 1997
- Three different types of Patterns:
  - Constant (C): 5 5 5 5 5 5 5 5 5 5 …
  - Stride (S): 1 2 3 4 5 6 7 8 9 …
  - Non-Stride (NS): 28 13 99 107 23 456 …
- Combinations:
  - Repeated Stride (RS): 1 2 3 1 2 3 1 2 3 1 2 3
  - Repeadted Non-Stride (RNS): 1 -13 -99 7 1 -13 -99 7

Computational Predictors

- Last Value Predictors
  - Predict that instruction will produce same value as last time
  - Requires some form of hysteresis. Two subtle alternatives:
    - » Saturating counter incremented/decremented on success/failure replace when the count is below threshold
    - » Keep old value until new value seen frequently enough
  - Second version predicts a constant when appears temporarily constant
- Stride Predictors
  - Predict next value by adding the sum of most recent value to difference of two most recent values:
    - » If v_{n-1} and v_{n-2} are the two most recent values, then predict next value will be: $v_{n-1} + (v_{n-1} - v_{n-2})$
    - » The value $(v_{n-1} - v_{n-2})$ is called the "stride"
  - Important variations in hysteresis:
    - » Change stride only if saturating counter falls below threshold
    - » Or "two-delta" method. Two strides maintained.
      - First (S1) always updated by difference between two most recent values
      - Other (S2) used for computing predictions
      - When S1 seen twice in a row, then S1⇒S2
- More complex predictors:
  - Multiple strides for nested loops
  - Complex computations for complex loops (polynomials, etc!)

Context Based Predictors

- Context Based Predictor
  - Relies on Tables to do trick
  - Classified according to the order: an "n-th" order model takes last n values and uses this to produce prediction
    - » So – 0th order predictor will be entirely frequency based
- Consider sequence: a a a b c a a a b c a a a
  - Next value is?

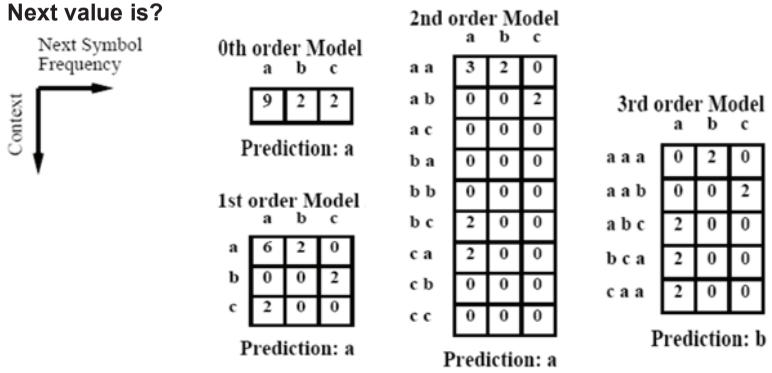

- "Blending": Use prediction of highest order available

Which is better?

- Stride-based:
  - Learns faster
  - less state
  - Much cheaper in terms of hardware!
  - runs into errors for any pattern that is not an infinite stride
- Context-based:
  - Much longer to train
  - Performs perfectly once trained
  - Much more expensive hardware

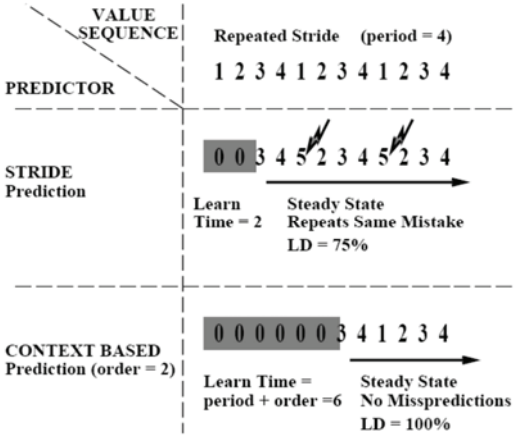

How predictable are data items?

- Assumptions – looking for limits
  - Prediction done with no table aliasing (every instruction has own set of tables/strides/etc.
  - Only instructions that write into registers are measured
    - » Excludes stores, branches, jumps, etc
- Overall Predictability:
  - L = Last Value
  - S = Stride (delta-2)
  - FCMx = Order x context based predictor

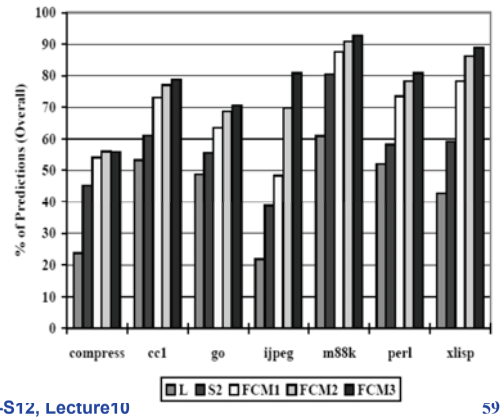

Correlation of Predicted Sets

- Way to interpret:
  - l = last val
  - s = stride
  - f = fcm3
- Combinations:
  - ls = both l and s
  - Etc.
- Conclusion?
  - Only 18% not predicted correctly by any model
  - About 40% captured by all predictors
  - A significant fraction (over 20%) only captured by fcm
  - Stride does well!
    - » Over 60% of correct predictions captured
  - Last-Value seems to have very little added value

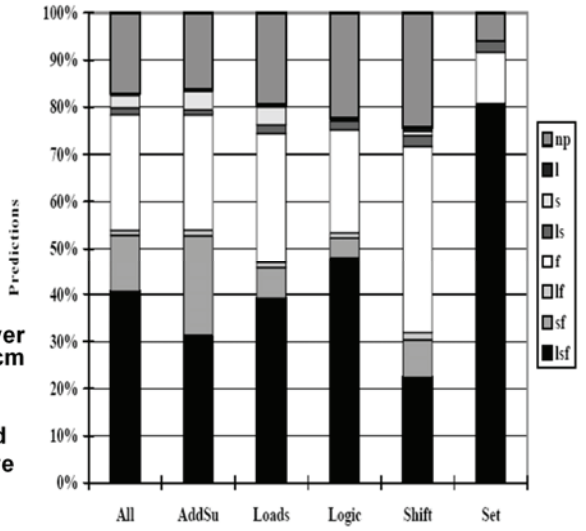

Number of unique values

- Data Observations:
  - Many static instructions (>50%) generate only one value
  - Majority of static instructions (>90%) generate fewer than 64 values
  - Majority of dynamic instructions (>50%) correspond to static insts that generate fewer than 64 values
  - Over 90% of dynamic instructions correspond to static insts that generate fewer than 4096 unique values
- Suggests that a relatively small number of values would be required for actual context prediction

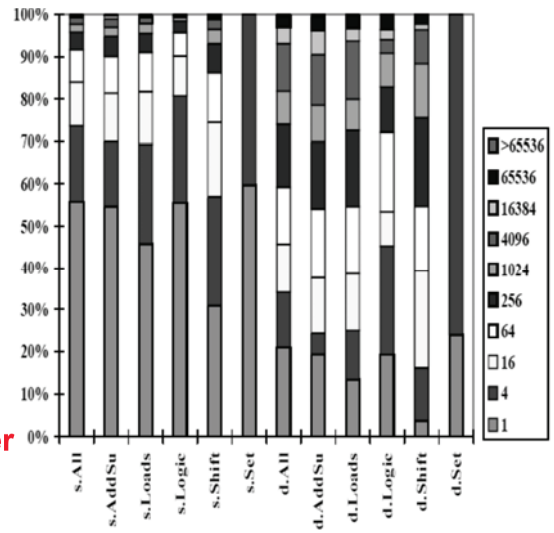

General Idea: Confidence Prediction

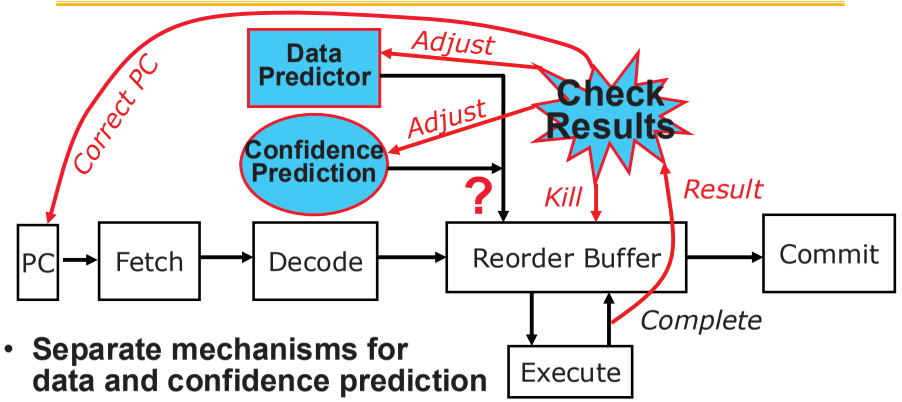

- Separate mechanisms for data and confidence prediction
  - Data predictor keeps track of values via multiple mechanisms
  - Confidence predictor tracks history of correctness (good/bad)
- Confidence prediction options:
  - Saturating counter
  - History register (like branch prediction)

Discussion of papers: The Alpha 21264 Microprocessor

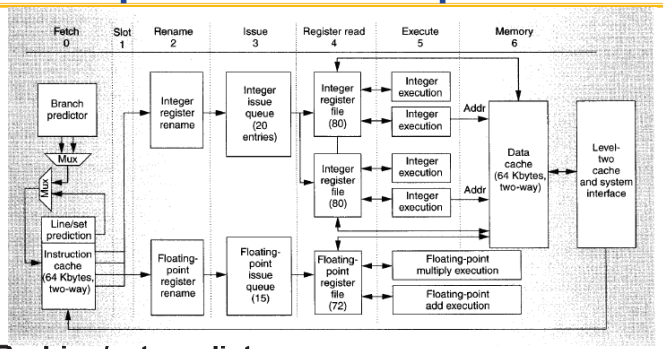

- BTB⇒Line/set predictor
  - Trained by branch predictor (Tournament predictor)
- Renaming: 80 integer registers, 72 floating-point registers
- Clustered architecture for integer ops (2 clusters)
- Speculative Loads:
  - Dependency speculation
  - Cache-miss speculation

Conclusion (1 of 2)

- Prediction works because….
  - Programs have patterns
  - Just have to figure out what they are
  - Basic Assumption: Future can be predicted from past!
- Correlation: Recently executed branches correlated with next branch.
  - Either different branches (GA)
  - Or different executions of same branches (PA).
- Two-Level Branch Prediction
  - Uses complex history (either global or local) to predict next branch
  - Two tables: a history table and a pattern table
  - Global Predictors: GAg, GAs, GShare
  - Local Predictors: PAg, Pap
- Aliasing: can be good *or* bad

Conclusion (2 of 2)

- Dependence Prediction: Try to predict whether load depends on stores before addresses are known
  - Store set: Set of stores that have had dependencies with load in past
- Last Value Prediction
  - Predict that value of load will be similar (same?) as previous value
  - Works better than one might expect
- Computational Based Predictors
  - Try to construct prediction based on some actual computation
  - Last Value is trivial Prediction
  - Stride Based Prediction is slightly more complex
- Context Based Predictors
  - Table Driven
  - When see given sequence, repeat what was seen last time
  - Can reproduce complex patterns